Thu, 29 Sep 2005, 9:04:41 AM EST

Edit an existing query or compose a new query in the Search Query Display.

Select a search number (#) to:

- Add a query to the Search Query Display
- Combine search queries using AND, OR, or NOT
- Delete a search
- Run a search

Search Query Display

Recent Search Queries

#1	(~~virtual machine~~ and thread and servlet and (spawn or fork or load) and available)<in>pdfdata
#2	(virtual machine server and thread and servlet and (spawn or fork or load) and available)<in>pdfdata
#3	(virtual machine server and thread and (spawn or fork or load) and available) <in>pdfdata
#4	(virtual machine server and thread and available)<in>pdfdata
#5	(virtual machine server and available)<in>pdfdata
#6	(virtual machine server)<in>pdfdata
#7	(virtual machines<in>de)
#8	(virtual machines loader <in>de)
#9	(virtual machine manager and available)<in>pdfdata
#10	(virtual machine manager and available)<in>pdfdata

Clear Session History

Indexed by Inspec

© Copyright 2005 IE

ACM PORTAL
USPTO
Subscribe (Full Service) Register (Limited Service, Free) Login
Search: The ACM Digital Library The Guide
"virtual machine hypervisor" and thread and available
SEARCH
THE ACM DIGITAL LIBRARY
Feedback Report a problem Satisfaction survey

Terms used **virtual machine hypervisor** and **thread** and **available**

Found **7,992** of **161,645**

Sort results by: relevance
Display results: expanded form
Save results to a Binder
Search Tips
Open results in a new window

Try an Advanced Search
Try this search in The ACM Guide

Results 1 - 20 of 200 Result page: **1** 2 3 4 5 6 7 8 9 10 next
Best 200 shown Relevance scale

1 Extended architecture and Hypervisor performance
Carl J. Young
March 1973 **Proceedings of the workshop on virtual computer systems**
Full text available: pdf(235.44 KB) Additional Information: full citation, abstract, citings, index terms

This paper is a brief summary of the impact that architecture extensions to hardware and software have upon the design and performance of software Hypervisors that are intended to provide the extended function in a virtual machine environment.

2 Asynchronous software thread integration for efficient software
Nagendra J. Kumar, Siddhartha Shivshankar, Alexander G. Dean
June 2004 **ACM SIGPLAN Notices , Proceedings of the 2004 ACM SIGPLAN/SIGBED conference on Languages, compilers, and tools for embedded systems**, Volume 39 Issue 7
Full text available: pdf(229.95 KB) Additional Information: full citation, abstract, references, citings, index terms

The overhead of context-switching limits efficient scheduling of multiple concurrent threads on a uniprocessor when real-time requirements exist. Existing software thread integration (STI) methods reduce context switches, but only provide synchronous thread progress within integrated functions. For the remaining, non-integrated portions of the secondary threads to run and avoid starvation, the primary thread must have adequate amounts of coarse-grain idle time (longer than two context-switches). ...

Keywords: J1850, asynchronous software thread integration, fine-grain concurrency, hardware to software migration, software-implemented communication protocol controllers

3 Compiler analysis and optimization: Providing time- and space- efficient procedure calls for asynchronous software thread integration
Vasanth Asokan, Alexander G. Dean
September 2004 **Proceedings of the 2004 international conference on Compilers, architecture, and synthesis for embedded systems**
Full text available: pdf(289.56 KB) Additional Information: full citation, abstract, references, index terms

Asynchronous Software Thread Integration (ASTI) provides fine-grain concurrency in real-time threads by statically scheduling (integrating) code from primary threads into secondary threads, reducing the context switching needed and allowing recovery of fine-grain idle time.

Unlike STI, ASTI allows asynchronous thread progress.Current ASTI techniques do not support procedure calls in the secondary thread because they lead to timing conflicts during static scheduling. ASTI requires knowing the sec ...

Keywords: asynchronous software thread integration, fine-grain concurrency, hardware to software migration, software-implemented communication protocol controllers

4 Converting thread-level parallelism to instruction-level parallelism via simultaneous multithreading

Jack L. Lo, Joel S. Emer, Henry M. Levy, Rebecca L. Stamm, Dean M. Tullsen, S. J. Eggers
August 1997 **ACM Transactions on Computer Systems (TOCS)**, Volume 15 Issue 3

Full text available: pdf(526.39 KB) Additional Information: full citation, abstract, references, citings, index terms, review

To achieve high performance, contemporary computer systems rely on two forms of parallelism: instruction-level parallelism (ILP) and thread-level parallelism (TLP). Wide-issue super-scalar processors exploit ILP by executing multiple instructions from a single program in a single cycle. Multiprocessors (MP) exploit TLP by executing different threads in parallel on different processors. Unfortunately, both parallel processing styles statically partition processor resources, thus preventing t ...

Keywords: cache interference, instruction-level parallelism, multiprocessors, multithreading, simultaneous multithreading, thread-level parallelism

5 Another C Threads library

G. Cattaneo, G. Di Giore, M. Ruotolo
December 1992 **ACM SIGPLAN Notices**, Volume 27 Issue 12

Full text available: pdf(674.39 KB) Additional Information: full citation, abstract, index terms

In this paper we present the main results achieved with the implementation of an efficient multithread library, hereinafter called DIA_Thread library, built on the top of a shared memory symmetric multiprocessor environment. The idea which led us to this project was the lack of a *minimal* common layer which, as far the multitasking concerns, should help us to make our applications portable among different architectures and/or operating systems. In fact, there are more and more multiprocess ...

6 Kernel-level scheduling for the nano-threads programming model

Eleftherios D. Polychronopoulos, Xavier Martorell, Dimitrios S. Nikolopoulos, Jesus Labarta, Theodore S. Papatheodorou, Nacho Navarro
July 1998 **Proceedings of the 12th international conference on Supercomputing**

Full text available: pdf(1.20 MB) Additional Information: full citation, references, citings, index terms

7 Relational profiling: enabling thread-level parallelism in virtual machines

Timothy Heil, James E. Smith
December 2000 **Proceedings of the 33rd annual ACM/IEEE international symposium on Microarchitecture**

Full text available: pdf(237.19 KB) ps(1.61 MB) Additional Information: full citation, references, citings, index terms
Publisher Site

8 Balancing register allocation across threads for a multithreaded network processor
Xiaotong Zhuang, Santosh Pande
June 2004 **ACM SIGPLAN Notices , Proceedings of the ACM SIGPLAN 2004 conference on Programming language design and implementation**, Volume 39 Issue 6
Full text available: pdf(429.85 KB) Additional Information: full citation, abstract, references, index terms

Modern network processors employ multi-threading to allow concurrency amongst multiple packet processing tasks. We studied the properties of applications running on the network processors and observed that their imbalanced register requirements across different threads at different program points could lead to poor performance. Many times application needs demand some threads to be more performance critical than others and thus by controlling the register allocation across threads one could impa ...

Keywords: multithreaded processor, network processor, register allocation

9 The Vector-Thread Architecture
Ronny Krashinsky, Christopher Batten, Mark Hampton, Steve Gerding, Brian Pharris, Jared Casper, Krste Asanovic
March 2004 **ACM SIGARCH Computer Architecture News , Proceedings of the 31st annual international symposium on Computer architecture ISCA '04**, Volume 32 Issue 2
Full text available: pdf(317.13 KB) Additional Information: full citation, abstract

The vector-thread (VT) architectural paradigm unifies the vectorand multithreaded compute models. The VT abstraction providesthe programmer with a control processor and a vector of virtualprocessors (VPs). The control processor can use vector-fetch commandsto broadcast instructions to all the VPs or each VP can usethread-fetches to direct its own control flow. A seamless intermixingof the vector and threaded control mechanisms allows a VT architectureto flexibly and compactly encode application ...

10 Enhancing software reliability with speculative threads
Jeffrey Oplinger, Monica S. Lam
October 2002 **Proceedings of the 10th international conference on Architectural support for programming languages and operating systems**, Volume 37 , 36 , 30 Issue 10 , 5 , 5
Full text available: pdf(1.47 MB) Additional Information: full citation, abstract, references, citings

This paper advocates the use of a monitor-and-recover programming paradigm to enhance the reliability of software, and proposes an architectural design that allows software and hardware to cooperate in making this paradigm more efficient and easier to program.We propose that programmers write monitoring functions assuming simple sequential execution semantics. Our architecture speeds up the computation by executing the monitoring functions speculatively in parallel with the main computation. For

11 Implementing Ada 9X features using POSIX Threads: design issues
E. W. Giering, Frank Mueller, T. P. Baker
October 1993 **Proceedings of the conference on TRI-Ada '93**
Full text available: pdf(1.49 MB) Additional Information: full citation, references, citings, index terms

12 Architecture and implementation of Web sites: A multi-threaded PIPELINED Web server architecture for SMP/SoC machines
Gyu Sang Choi, Jin-Ha Kim, Deniz Ersoz, Chita R. Das
May 2005 **Proceedings of the 14th international conference on World Wide Web**

Full text available: pdf(259.39 KB) Additional Information: full citation, abstract, references, index terms

Design of high performance Web servers has become a recent research thrust to meet the increasing demand of network-based services. In this paper, we propose a new Web server architecture, called multi-threaded PIPELINED Web server, suitable for Symmetric Multi-Processor (SMP) or System-on-Chip (SoC) architectures. The proposed PIPELINED model consists of multiple thread pools, where each thread pool consists of five basic threads and two helper threads. The main advantages of the proposed model ...

Keywords: asynchronous multi-process event-driven, multi-process, multi-thread, single event-driven process, symmetric multi-processor, system-on-chip

13 A threaded/flow approach to reconfigurable distributed systems and service primitives architectures

L. F. Ludwig

August 1987 **ACM SIGCOMM Computer Communication Review , Proceedings of the ACM workshop on Frontiers in computer communications technology**, Volume 17 Issue 5

Full text available: pdf(1.19 MB) Additional Information: full citation, abstract, references, index terms

This paper discusses a methodology for managing the assembly, control, and disassembly of large numbers of independent small-scale configurations within large-scale reconfigurable distributed systems. The approach is targeted at service primitives architectures for enhanced telecommunications networks, but can apply to more general settings such as multi-tasking supercomputers and network operations systems.* Study of the methods presented here was a key motivation in f ...

14 The STAMPede approach to thread-level speculation

J. Gregory Steffan, Christopher Colohan, Antonia Zhai, Todd C. Mowry

August 2005 **ACM Transactions on Computer Systems (TOCS)**, Volume 23 Issue 3

Full text available: pdf(1.72 MB) Additional Information: full citation, abstract, references, index terms

Multithreaded processor architectures are becoming increasingly commonplace: many current and upcoming designs support chip multiprocessing, simultaneous multithreading, or both. While it is relatively straightforward to use these architectures to improve the throughput of a multithreaded or multiprogrammed workload, the real challenge is how to easily create *parallel software* to allow single programs to effectively exploit all of this raw performance potential. One promising technique fo ...

Keywords: Thread-level speculation, automatic parallelization, cache coherence, chip-multiprocessing

15 Implementing an efficient vector instruction set in a chip multi-processor using micro-threaded pipelines

Chris Jesshope

January 2001 **Australian Computer Science Communications , Proceedings of the 6th Australasian conference on Computer systems architecture ACSAC '01**, Volume 23 Issue 4

Full text available: pdf(911.26 KB) Publisher Site Additional Information: full citation, abstract, references, citings

This paper looks at a combination of two techniques, one of which, using a vector instruction set, has a long history dating back to pipelined vector supercomputers, such as the Cray 1 and its successors. The other technique, multi-threading, is also well understood. The novel approach proposed in this paper combines both vertical and horizontal micro-threading with

vector instruction descriptors. It will be shown that a family of threads can represent a vector instruction with dependencies betw ...

16 On the multi-threaded computation of integral polynomial greatest common divisors
Wolfgang Küchlin
June 1991 **Proceedings of the 1991 international symposium on Symbolic and algebraic computation**
Full text available: pdf(1.04 MB) Additional Information: full citation, references, citings, index terms

17 Providing soft real-time QoS guarantees for Java threads
James C. Pang, Gholamali C. Shoja, Eric G. Manning
June 2001 **Proceedings of the 2001 joint ACM-ISCOPE conference on Java Grande**
Full text available: pdf(853.96 KB) Additional Information: full citation, abstract, references, index terms

The Java platform has many characteristics that make it very desirable for integrated continuous media processing. Unfortunately, it lacks the necessary CPU resource management facility to support quality of service guarantees for soft real-time multimedia tasks. In this paper, we present our new Java Virtual Machine, Q-JVM, which brings CPU resource management to the Java platform. Q-JVM is based on Sun's JVM version 1.1.5. It implements an enhanced version or the MTR-LS algorithm in its thr ...

18 Explicit multi-threading (XMT) bridging models for instruction parallelism (extended abstract)
Uzi Vishkin, Shlomit Dascal, Efraim Berkovich, Joseph Nuzman
June 1998 **Proceedings of the tenth annual ACM symposium on Parallel algorithms and architectures**
Full text available: pdf(1.71 MB) Additional Information: full citation, references, citings, index terms

19 KDB: a multi-threaded debugger for multi-threaded applications
Peter A. Buhr, Martin Karsten, Jun Shih
January 1996 **Proceedings of the SIGMETRICS symposium on Parallel and distributed tools**
Full text available: pdf(991.10 KB) Additional Information: full citation, references, citings, index terms

20 Object and native code thread mobility among heterogeneous computers (includes sources)
B. Steensgaard, E. Jul
December 1995 **ACM SIGOPS Operating Systems Review , Proceedings of the fifteenth ACM symposium on Operating systems principles**, Volume 29 Issue 5
Full text available: pdf(1.50 MB) Additional Information: full citation, references, citings, index terms

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2005 ACM, Inc.
Terms of Usage Privacy Policy Code of Ethics Contact Us

Useful downloads: Adobe Acrobat QuickTime Windows Media Player Real Player

Searched all groups

Results 1 - 10 of 57,100 for virtual machine thread (

Sorted by relevance Sort by date

java virtual machine stack

... 3.5.2 Java **Virtual Machine** Stacks Each Java **virtual machine thread** has a private
Java **virtual machine** stack, created at the same time as the **thread**. ...
comp.lang.java.programmer - May 13 2003, 5:20 pm by Doug Pardee

Virtual Machine in XP

... (see new post on previous discussion **thread**) I cannot ... to get something as simple
as the VM, please feel ... open major school labs with Linux **machines** instead of ...
microsoft.public.windowsxp.general - Aug 2 2002, 5:03 am by Mike Brannigan [MS] - 12 messages - 8 authors

OutOfMemory Virtual Machine Inconsistency/Strangeness

... unable to create new native **thread** at java.lang.**Thread**.start(Native Method) I can
execute my program several times with the **virtual machine** opening, running ...

comp.lang.java.programmer - Mar 18 2003, 5:13 pm by Iain Toft

java virtual machine

... After that, the utility should work properly." -- Please respond in the same **thread**. ...
microsoft.com> typed: I need to download and install java **virtual machine**
microsoft.public.windowsupdate - Oct 27 2004, 6:05 am by Robert Aldwinckle - 13 messages - 8 authors

start up failure after snapshot failed

... Release.4.0.0 Nov 19 23:04:08: mks| VTHREAD declare **thread** 0 "vmx ... tooOld] One or
more of the **virtual** disks configured with this **virtual machine** were created by ...

vmware.for-linux.general - Nov 19 2003, 7:32 pm by kbailey - 4 messages - 1 author

Sponsored Links

VMware **Virtual** Machines
Download Features, Product Specs & a Free Workstation 5 Trial Today
www.vmware.com

Discount **Thread** & More
Wholesale pricing with no minimums Embroidery and decorative threads
www.threadart.com

Pipe Threading Machines
Portable Pipe Threaders by Wheeler-Rex. Contractor Discounts Available
www.BuildMartStore.com

Threading **Machine**
Quality new and used items. Find threading **machine** now!
www.ebay.com

Threading Equipment Co.
Huge Inventory of Die Heads, Dies, Chasers, Tapping Tools & Gages
threadingequipmentco.com

Virtual Machine
Free **Virtual Machine** info from the experts at the Tech Encyclopedia.
www.Tech-Encyclopedia.com

Find Threading **Machine**
We have what you're looking for. Threading **Machine** & much more!
www.eWoss.com

See your message here...

Distributed **virtual machine**
... and there have been many attempts to implement **virtual** shared memory ... well-designed
parallel databases, the migration of **threads** between **machines**, I/O ...
comp.parallel - Oct 9 2004, 5:36 pm by Nick Maclaren - 2 messages - 2 authors

Java **virtual machine** not compatable w/ xp
... ago, before posting this message, I used the link provided by Mark Salloway on June
12, 2002, in the "java **Virtual Machine**" **thread**: http://makeashorterlink.com ...
microsoft.public.windowsxp.games - Aug 13 2002, 5:13 pm by Chris H. - 9 messages - 6 authors

Threads Design. A Win32 perspective.
... t run a even a single - the one and only - **thread** without these ... seraration of the
VMstate from the interpreter state, VM = **Virtual machine** = interpreter These ...
perl.perl6.internals - Jan 3 2004, 8:00 pm by Leopold Toetsch - 49 messages - 12 authors

ANNOUNCE: JPE, the Java-Python Extension - first beta release
... Each **thread** is attached to zero or one **thread** environment from each **virtual machine**; each **thread** environment is attached to one **thread**. ...
comp.lang.python - Apr 29 2001, 8:01 pm by Frederic Giacometti - 6 messages - 2 authors

Thread in a Java **Virtual Machine** : Any Doc ?
I found the Java **Virtual Machine** specs, and the .class file format, but not a word on how multitasking ins handled by a **Virtual** Java **Machine**, Especialy, what ...
comp.lang.java - Jun 12 1996, 2:44 am by Remi Arnaud - 1 message - 1 author

New! Get the latest messages on **virtual machine thread** emailed to you with Google Alerts.

Gooooooooogle ▶
Result Page: 1 2 3 4 5 6 7 8 9 10 Next